

The Human Interface Group

Space Planning Project Technical Report

Leonardo Project

Lewis Knapp, team leader

Penny Bauersfeld

Michael Chen

Stephanie Houde

Laurie Vertelney

Prepared by:

Penny Bauersfeld

Human Interface Group

Apple Computer, Inc.

November 30, 1990

with contributions from Michael Chen, Lewis Knapp and Laurie Vertelney

Technical Report 90-15

TABLE OF CONTENTS

Introduction	2
Background	3
Development	5
Project Overview	5
User Observation	5
Interface Approaches	6
Scenario-Based Design	7
Prototyping	8
Working Prototypes	10
Handles: Intuitive 3D Manipulation with Existing Hardware	10
Thumbwheel Chair: Object Manipulation with Alternate Device	12
People View: Using People to Manipulate Views	14
Thumbwheel View: View Manipulation with Alternate Device	15
Integrated Application Prototype	18
Integrated Functionality and Constraints	18
Available Manipulation Methods	19
World and View Manipulations	19
User Testing and Iteration	20
Incorporation of Localized Sound	21
3D Sound	22
Findings	23
Interaction Based on Users' Expectations	23
Implications for Developers	23
Manipulating People v. Objects	24
Flexible Interchange of Methods	24
Role of User Testing	25
3D Prototyping Isn't Easy!	25
Issues	26
Input Device	26
Device mappings and user customizability	26
Integrated Interface	26
Role of sound	27
Object Constraints and Default Behaviors	27
Conclusions	28

INTRODUCTION

The Leonardo project team was formed to explore interfaces for 3D interaction and design. While a number of individual efforts have been undertaken as part of the Leonardo project, the team members came together in the summer of 1990 to focus on the problem of a direct manipulation interface for 3D interaction. The Space Planning project encompasses this work. Space Planning, or specifically moving 3D objects (such as furniture) around in 3D space, concentrates on the manipulation of 3D objects rather than the formative design or modeling of 3D objects.

The approach for the Space Planning project was to focus on the role of the user, or people in general, when manipulating 3D objects and spaces. The target users for the Leonardo project are designers who are not necessarily trained in the use of computers or geometry and formalisms, but who could benefit from increased visualization and simulation which the computer might provide. To address the needs of users, the team based its work on observations of potential users completing 3D manipulation tasks. From earlier studies of work by Paul Zimmerman, a local architect, the Leonardo team had determined the underlying importance of including people in 3D space¹; it is the role or position of people in a 3D space that help to give the space scale, function and purpose. By beginning the Space Planning project with the observations of users completing furniture manipulation tasks with a physical model, and keeping the user and his role in mind throughout the prototype development phase of the project, we were able to maintain this focus on the role of people in 3D manipulation. This focus on people is evident in many of the prototypes developed and presented in this report.

The major goal of the Space Planning project was to explore possibilities for direct manipulation 3D interfaces, and to make suggestions for 3D interaction based on the work completed. We hoped to demonstrate 3D interactive capabilities on the Macintosh platform, a machine not necessarily dedicated to 3D graphics. The intent was to employ a standard hardware configuration, including a simple input device. This report summarizes the Space Planning project. The highlights from the summer work are presented, and the findings and related issues discussed.

¹ Further information about the work with Paul Zimmerman is available in the separate technical report [Architect at Work Technical Report](#).

BACKGROUND

Use of computers for 3D design has evolved over the past 25 years from demonstrations on specialized hardware in university research settings to widely available application programs which run on personal computers. This evolution has been shaped by the advances in graphics hardware and computing capacity as well as by the needs of the engineering disciplines which could cost-justify using 3D graphics to improve their products. 3D graphics hardware and software development have thus in the past been driven by customers in the aerospace, automotive and other manufacturing industries.

The evolution of 3D graphics followed automation of drafting (a set of conventions for drawing 3D objects on 2D media) as acceptance levels of computer methods became higher in these industries. Commercial products such as generalized mechanical modelers became available, displacing the software developed in-house to meet specific corporate requirements. Graphics hardware which was capable of real-time motion and highly interactive behavior remained expensive, motivating vendors to use less capable hardware to run the same programs. Because of the time and effort required to add functionality to a piece of design software while at the same time porting it to each new generation of hardware, the user interface to the software evolved slowly. Interfaces were based on metaphors of drafting, mechanical engineering and the geometry of three-dimensional objects, probably because they were for the most part designed by engineers for whom such metaphors seemed most logical. The visual quality of the output produced by these systems has typically been much closer to line-oriented drafting than to "artist's renderings," because of the precision of the drafting abstraction which makes it easily understandable by a computer. Visual quality of the images has also been hindered by the higher cost of advanced color graphics hardware and software, and the computation entailed by complex renderings.

The result is that many of today's 3D design and engineering automation products require training in engineering-oriented dialogs in order to input, modify and view 3D data using a computer. Conventions for establishing views of an object require knowledge of the coordinate system in which an object is defined, precise numeric values for locations and dimensions and familiarity with the 2D geometric construction methods of drafting. In general, more consideration has been given to consistency of a software interface over time and across different hardware machines to preserve the customer base's training investment than to provide human engineered interfaces which meet users' expectations and are easy to use.

Even the new Macintosh-based 3D design programs which are available carry many of the traditional numeric-oriented approaches to view control and model creation, requiring the user to understand a common abstraction or a new, stylized scheme. Admittedly, some of the interfaces have become less numeric by adopting sliders and thermometers which remove some of the detail, but which rely on an indirect control mechanism which still require users to understand abstract software models of the program.

With the Space Planning project, we set out to build a 3D direct manipulation interface to illustrate a process of interaction which is simple enough to be approachable by a wide audience of potential users. We realize that such an interface may seem deliberately oversimplified to the technically astute computer technologists who traditionally have built such applications, but we hoped to address a user population beyond this minority of 3D users. We wanted to build a prototype to provide a basis for further discussions of how similar capabilities could be implemented consistently across a family of applications, and to show that simple concepts like direct manipulation, which make the Mac easy to use, generalize directly to 3D interactions. The prototype is intended to evoke discussion of other ways manipulation might be made available, not to be a complete and definitive solution to everyone's problems.

Our approaches were based on previous work and general assumptions about what we believed the non-engineering user would like to use. We tried to use the expectations which people already have, based on how real world things work. We tried to understand what cues helped interaction with 3D objects and what it would take to provide these cues in the interface. We wanted to push in the direction of using the cues which make interaction with 3D world objects imitate the things which we see and feel (and hear) when we do real 3D tasks.

Our choice of the space planning theme was frequently associated with the previous two*Something-D* project in HIG. We believe that while both two*Something-D* and the Space Planning projects had similar application areas, they were fundamentally different because we concentrated on the direct manipulation aspect of the interaction whereas two*Something-D* had emphasized the use of voice and delegation to perform furniture placement tasks. While two*Something-D* was about the combination of a voice delegation and direct manipulation interface, and might have chosen a different application area with which to demonstrate its principles, the Space Planning project inherently depended on the task of 3D furniture manipulation.

DEVELOPMENT

Project Overview

The basic task of Space Planning was chosen to showcase direct manipulation of 3D objects, specifically furniture objects in a living room. This project differs from other 3D work underway at Apple because of this focus on direct manipulation of 3D objects, rather than on modeling or rendering 3D objects. By concentrating on the manipulation of objects, without the complications of 3D modeling, we hoped to introduce elemental ways of interacting with 3D objects. We hoped that the theme of furniture arrangement is one with which many people are familiar and which they can envision themselves doing. The interactions presented in the Space Planning prototypes include sliding 3D objects on a plane (the floor), rotating 3D objects (around their own central axes) and lifting 3D objects off the plane. Other 3D interface components, such as manipulating the view of the 3D space, are also introduced.

User Observation

To begin the project, we conducted an initial user observation to try to understand the types of manipulations people might anticipate in direct manipulation interaction with 3D objects. We find that such informal observations allow us to then design an interface based on user expectations. We hoped to be able to extract findings from the physical furniture manipulation model which would help us to design an on-line interactive furniture manipulation interface. Six Apple employees of varying skills were asked to interact with a foam core model of a furnished room. The room model contained a desk, a table, two chairs and two pictures hanging from the walls. Users completed a series of tasks which encouraged them to interact with different pieces of furniture in different ways, attempting to understand which types of interactions were used for which tasks.

The basic findings of the observation were: 1) people naturally understood interaction with the 3D objects and recognized the need for 3D manipulation in 3D space; 2) each person manipulated the furniture somewhat differently, leading us to believe that any method of 3D interaction must encourage individuality or flexibility in the functions it provides; 3) most of the manipulations could be categorized as sliding an object along a plane, lifting an object, or re-orienting (rotating) an object in space; and 4) users had expectations about the way objects would be constrained, based on the object being manipulated.²

² Findings are documented in the paper Space Planning: Leonardo 3D Manipulation Project / Findings from Preliminary User Observations. Video tapes are available.

The user observations influenced our work in a number of ways. While they indicated that each user went about the 3D manipulation tasks somewhat individually, and that we would thus have to provide a flexible means of interaction, we were encouraged to see that the same basic vocabulary for object manipulation was used by all subjects. We concluded that a tightly modal, single-path driven interactive approach would not be appropriate, as is typically found in existing 3D packages. Instead, we decided to build a number of prototypes which could explore alternate methods of interaction. Eventually we hoped to combine interactive qualities from the various prototypes to come up with a flexible interface to directly manipulate 3D objects. The interactive tasks we would provide would include the slide, rotate and lift functions as identified in the user observations.

Interface Approaches

The user observations led us to examine the ways we might approach interfaces for 3D manipulation. First, we observed that each user felt he had a definite role in the space planning activities, often speaking about the furniture as if he was in the space with it. We realized that the role of the user, and of people in general, is important in a 3D environment, probably much more so than in a 2D world. Users relate the 3D space represented on the screen to a 3D space that they might occupy, and readily see themselves *inside* the space with which they are working. Users could not have such a relationship with a 2D space such as is presented in a word processing application.

The notion that users see themselves in the 3D scene led us to believe that placing people in the 3D scene might be a good metaphor for specifying points of view of the space. The use of people in the scene also provides important scale information about the other objects in the room, and further enhances the connection between the user and the space. This approach that most users took, linking themselves to the 3D space, led us to believe that in general, the tasks involving object manipulation might be approached somewhat differently than those involving the manipulation of people. We recognized that both types of manipulation are desirable in a 3D space. Because users see themselves as having an active role in the arrangement of furniture, they see the need for people and control of these people in the 3D space they are arranging. The observations showed us, however, that while users actively directly manipulated the pieces of furniture in the room, they talked about more passive interactions with people who might occupy the space.

This led us to discuss the ways in which we might want to approach people in the

3D space. It was clear that we would need to support different views of the space for different people, or at least allow the user to manipulate a single view to correspond to different characteristics of different people. We hoped to take advantage of the computer's power to simulate various conditions to be able to view the 3D space from different view points. For instance, we could imagine viewing the space first as ourselves, then asking the computer to simulate the view from the same position if the viewer were a child, or handicapped person, or if the viewer were taller. We discussed the possibilities of incorporating these human scale factors into the interface as a more intuitive interface for changing one's point of view in a scene. This is in contrast to the more abstract method of specifying the camera's location, *look at* point and *up-vector* that many 3D packages employ to change one's view.

We wanted to continue development of the virtual controllers in different contexts because of their demonstrated appeal and measured performance in prior prototypes. For example, we wanted to show how virtual controllers can be applied to an environment where objects have built-in constraints.

Finally, we realized that while many opportunities existed for both object and people manipulation, it would be most important not to present too many confusing methods of interaction to the user at once. We saw the need to maintain consistency among the approaches we might recommend. Early prototypes would need to explore the many possibilities for interaction, but as we began to combine prototypes we must keep in mind the need for consistency and clarity among the approaches chosen.

Scenario-Based Design

To help identify tasks for 3D manipulation, we designed a user scenario for furniture arrangement based on a real-life situation. The base scenario chosen was the rearrangement of furniture in a living room to make room for a new television stand. In the process of setting up the new stand, the user needed to move some of the older furniture aside, rearrange the furniture, enter the space to view the new effects, make changes based on this view, and then view the space again. Twenty key frames were identified for the scenario, and specific manipulations were identified to get from one key frame to the next. The transitions between consecutive key frames were used as the starting points for the ensuing prototypes. In this way, real user tasks drove the development of the prototypes, and the motivation for interaction was established.³

³The complete scenario is documented in the paper Space Planning: Leonardo 3D Manipulation Project / User Scenario.




Figure 1: The living room from the user scenario. The various items in the room are manipulated to support the identified 3D interaction tasks.

Prototyping

We began prototyping the scenario key frames via various approaches with a single generalized goal: to use the tools that were available to us to try to present various possibilities for 3D interaction. We wanted to visualize our interface ideas as quickly as possible, without having to invest too much time building supporting tools. We chose MacroMind Director as a primary prototyping tool because we believed it to be the most appropriate tool for our needs, supporting color and the importation of images modeled in Swivel 3D. We chose Director also for its support of imported sound clips, and because of its XObject interface which would allow us to hook in alternate interface elements that might not otherwise be supported. We hoped to explore interface prototypes as quickly as possible in Director, then transfer the concepts we had developed to an integrated application written in MacApp. Michael Chen developed this MacApp interface in parallel, as approaches tested in the earlier Director prototypes were made clearer.

We made some initial decisions about the nature of the interface. We decided early on that the input device used for the interface would be important in driving the interactive experience. We did not want to present a complex or unnatural method of interaction to the user, so we chose to use only simple input devices, limiting our prototypes to use either the standard Macintosh mouse or the thumbwheel mouse. The thumbwheel mouse (see Figure 2 below), developed by Dan Venolia in the Human Interface group, is a standard mouse with a roller wheel attached to the side and accessible by the user's thumb. This wheel allows the user to control an additional degree of freedom, which can be mapped to alternate input axes in the interface. We chose the thumbwheel mouse because

we wanted to prototype with a device that afforded an additional degree of freedom of manipulation. We chose not to give the user full six-degree of freedom manipulation because we believed that available six-degree of freedom devices, such as the polhemus-based cue ball or the Dataglove, can be awkward and difficult to control. Because we wanted object selection to be straightforward, we decided to present the user with a 2D cursor, mapping to the two dimensions of the screen rather than the three dimensions of the space depicted, since we believed this would give the user greater control over the environment. The user would select an object by pointing to the object on the surface of the screen, rather than having to locate the object in 3D space.

Finally, we hoped to provide interfaces that would run on a general purpose machine, not specifically dedicated to 3D graphics. Each of the prototypes was implemented on a Macintosh, with 8 meg of internal memory (necessary for the large Director prototypes). The scenario for interaction was developed by Penny Bauersfeld, Laurie Vertelney and Stephanie Houde. The living room used for the furniture manipulation tasks was designed by Laurie Vertelney and modeled by Stephanie Houde using Swivel 3D. Each prototyper used the Swivel models and the tweening function provided in Swivel to build animations for the various prototypes, and then brought them from Swivel into the Director environment. The integrated prototype by Michael Chen was programmed using MacApp and a specially designed custom 3D graphics package. Director drivers for the thumbwheel mouse were written by Dan Venolia and Erik Neumann and adapted for our use by Penny Bauersfeld. Voice clips recorded using MacRecorder and SoundEdit and sounds from various public domain CDs were imported into the Director prototypes using the existing sound functions.




Figure 2: A user using the Thumbwheel Mouse.

WORKING PROTOTYPES

The prototypes described in this section were built in MacroMind Director. All except the PeopleView prototype are available as interactive simulations. The PeopleView prototype is a stand-alone animation.

Handles

Intuitive 3D Manipulation with Existing Hardware

Stephanie Houde, a summer intern with the Leonardo team, set out to solve the problem of 3D interaction using the existing I/O interface, namely the Macintosh with a standard mouse. The key frame transition from the scenario she focused on was moving a chair across the living room to make room for the new position of the television stand. While we recognized that 3D interaction may be most intuitively accomplished with new I/O modalities, we wanted to explore more user-driven methods of interaction with existing hardware. Existing 3D applications present 3D manipulation in non-intuitive, model-based ways that we felt could be improved upon without additional devices. Stephanie's summer research focused on how people interact with objects, and she explored the ways in which objects might be directly manipulated by users in the most intuitive manner. The overall research project is referred to as Handles, for the interactive devices used to manipulate objects. A separate technical report, Hands On 3D Object Manipulation by Stephanie Houde is available for further details on the Handles project; a basic summary is given here.




Figure 3: A chair object from the Handles prototype is shown in its selected state, with a handle box. Any of the handles are now available for interaction. The handle at the top is used to lift the chair, the handles in the lower corners to rotate it.

The Handles project focused on the use of visual cues, in the form of handles to objects, to show how the user might interact with these objects. For instance, a chair might appear with handles for manipulation as shown above. The graphic

space. With existing two-dimensional input devices, the user can perform only one of these activities at a time, however with the introduction of the thumbwheel mouse the user can perform both simultaneously, as he would in the physical model. Preliminary findings based on casual use of the Thumbwheel Chair prototype show that it seems natural and appealing to users.




Figure 4: View of the screen from the Thumbwheel Mouse prototype. The chair in the center of the room is controlled with the thumbwheel mouse. The window in the upper right corner is used to indicate that the chair is currently selected.

One major issue with the existing Thumbwheel Chair prototype is that because the thumbwheel mouse only provides one additional axis for manipulation, it does not support all needed functionality for furniture manipulation. With the current mapping, the user cannot lift the chair off the floor. Additional development might be necessary to allow the user to customize the mapping of the input device to allow for the full range of potentially desired functions. The mapping could start out as is, with the mouse movement mapping to moving the chair and the thumbwheel roller movement mapping to rotating the chair, but might be changeable so that the user can choose an alternate mapping when desired to allow, for example, the roller wheel movement to map to lifting the chair off the floor. Such customization is left for further research.

handles indicate a number of things to the user simultaneously. They show where on the object the user could select to perform a manipulation. By their form, they indicate the type of interaction available to the user. For instance, if the user selects the handle atop the object the ensuing manipulation will be to lift the object. If a handle along the lower corner is selected, the user will be able to rotate the object. By selecting anywhere else on the object, the user will be able to translate or move the object, much as a standard select-and-drag interaction works in 2D applications. Handles thus allow the user to directly interact with the object without having to separately and awkwardly select a mode (such as 'lift' or 'rotate') for interaction, as is common in available 3D interactive applications such as Swivel 3D or Super 3D. Here, because the object being manipulated is a piece of furniture, human hands are used as the representation of handles, both to help give scale to the size of the object and to provide information about how the object can be manipulated. With other objects, such as 3D chemical molecule models, we might imagine alternate handle objects.

An important finding in the Handles research, found in earlier prototyping, is that a familiar form like the bounding box around the object is a necessary element of the interaction with 3D objects. An early informal study showed that, given a certain object like a wing-backed chair, different users had completely different approaches to manipulating that object. Some users would attempt to grab the chair by the arm, some by the back, some from beneath. This behavior follows also from the preliminary user observation with the physical furniture model, where different users each manipulated the furniture somewhat uniquely. In the 'real' world, users can grab 'real' 3D objects any way they choose, but in the computer modeled world such random handling would be very difficult to track, and the computer would not be able to predict what the user really wants to do when he clicks on part of an object and moves the mouse. By placing each object, as the chair above, in a 3D bounding box, the user no longer need make decisions about where to grab the particular object. All objects are treated similarly, as if they were located within a 3D rectangular box, and the handles are clearly indicative of the interactions that can occur with this box. Using the bounding box around the object helps to generalize the approach to 3D interaction.

The basic findings from informal user studies of the Handles interface is that the interaction seems natural and smooth. Users seem to understand from the start how to manipulate objects with handles and how to interact with the handles themselves. While formal testing is yet to be completed, we believe that overall, handles provide an excellent way to interact with 3D objects using a conventionally configured Macintosh.

Thumbwheel Chair

Object Manipulation with Alternate Device

In addition to exploring interfaces for existing hardware configurations, we believe it is important to research alternate interface devices for object manipulation. In the 'real' world, users move furniture objects in three dimensions, with up to six degrees of manipulation freedom. Even though the furniture objects in the Space Planning scenario exist in a 3D world, six degrees of freedom may not be necessary for smooth interaction. For example, a chair will normally only be slid on the plane of the floor or turned, movements which require only two or three axes of manipulation. With the existing Macintosh mouse, users are limited to two degrees of freedom. By providing interactivity with an alternate device such as the thumbwheel mouse introduced earlier, we can provide users with an additional simultaneous axis for manipulation. We hoped to use a device such as the thumbwheel mouse to provide more intuitive methods of interaction. We believed the thumbwheel mouse would be intuitive to use since it is an extension of a device with which users are already familiar, rather than a totally different device.

As noted in the *Prototyping* section above, the thumbwheel mouse provides an additional control through a roller wheel located at the position of the user's thumb. The user can roll the thumbwheel at the same time as he moves the mouse across the desktop, or can perform either task in isolation. The roller, because it is bidirectional, lends itself well to interaction of a continuous nature, such as sliding up and down a scroll bar⁴, or turning an object around an axis. Likewise, the roller might be used to access a third axis in 3D space, namely the z-axis, to complement the mouse's sliding action which accesses the x- and y-axes.

To explore using the thumbwheel mouse for object manipulation, Penny Bauersfeld built a prototype based on the same key frame transition from the scenario as in the Handles project, moving the chair from one side of the living room to the other, allowing the user to interact with a chair object. Using the same key frame transition sequence allowed us to have a basis for comparison of the two interactive methods. In the Thumbwheel Chair prototype, standard mouse movement maps to the chair sliding on the floor (as depicted in the 3D space). Movement of the roller maps to rotating the chair around its own central axis. The user can thus move and rotate the chair simultaneously by sliding the mouse and at the same time rolling the thumbwheel roller. This mapping was chosen to address a common approach to object manipulation indicated by the users in the preliminary user observations. Users of the foam core furniture often rotated furniture objects as they moved them from location to location in the 3D

⁴An early thumbwheel mouse prototype developed by Dan Venolia allowed the user to continuously scroll the currently selected window.

PeopleView

Using People to Manipulate Views

As we discovered from the informal user observations, users think about manipulating people differently from other inanimate objects. We believe the major differences stem from the fact that people are self-manipulating, that is, they can move themselves in space. In the real world, we rarely physically move other people, and they rarely move us. Instead, each person is responsible for moving himself. While we may direct another person to move himself, our involvement in moving others is typically delegatory, rather than through direct manipulation. Another factor in dealing with people in 3D spaces is that people are tied to the scene they view by the way they see the scene, directly through their eyes. A person's view of a scene is dependent on his placement in that space. These thoughts about the relationship of people to the scene led us to a number of ideas about using people in a scene to control the user's view of the scene.

Laurie Vertelney built a number of animated prototypes demonstrating some of the ways in which people might be manipulated by the user, and ways in which they can be used to control the view of the scene. She based her interactions on the scenario key frame transition sequences in which people were placed in the space and manipulated to see the effects on the view of the space. In these prototypes, she explores different methods of using the standard mouse to place people in a scene and adjust their position either standing or sitting in space. A number of interface elements, such as graphical menus, tool palettes and special effects, are used to illustrate potential interaction techniques. We believed that the Handles approach probably would not be appropriate for manipulating people, since the idea of 'man-handling' people to move them seemed somewhat offensive, or impolite at best. Instead, these people prototypes explored other interface possibilities with the standard mouse.

The People View prototype also explored the use of multiple views in a 3D interface. The different views provided alternate perspectives on the 3D scene. By providing multiple views in which to place and manipulate people, the prototype was able to explore alternate modalities. For example, in a smaller perspective view, the idea of a cone of vision to represent what the user might see in the larger view of the room was introduced. As the people in the scene were moved to provide alternate views, the cone of vision updated accordingly. With multiple views of the same scene, the user could maneuver whichever view seemed most natural for a given task. Multiple views also allow the user a more complete model of the 3D space.




Figure 5: A scene from the People View prototype. Two people, a male adult and a small child, are available for view manipulation.

The idea of using information about a specific person to control the characteristics of a person in the scene was also incorporated in these prototypes. We incorporated ideas from the Human Scale publications⁵ in some of the interactions, allowing users to specify whether the person was a man, woman or child, and whether the person was sitting or standing. We imagine that in a full scale interface, users would have access to height and size components of people. With the ability to specify these traits, the user could see what the scene might look like through the eyes of a man of certain height vs. the child of another height, or what it might look like through the eyes of a handicapped person. Likewise, furniture could be "sat in" by people of varying weights and sizes, to model potential comfort in an office or living room. These People View prototypes introduce these issues and the potential for further development.

Thumbwheel View

View Manipulation with Alternate Device

As with objects, we chose to explore the use of alternate input devices in prototyping people and view manipulation. Because the user's view of the scene is three-dimensional, users may benefit from input devices with additional axes of control to manipulate the view onto the space. Penny Bauersfeld introduced a

⁵Diffrient, Niels S., Tilley, Alvin R., and Bardagly, Joan C. Humanscale™1/2/3: A Portfolio of Information. Cambridge, MA: The MIT Press, Henry Dreyfuss Associates, designers, 1974.

number of prototypes which explore the use of the thumbwheel mouse for manipulating oneself (as the user) within the 3D space. As in the object manipulation prototypes with the thumbwheel mouse, the mapping of the device to the interaction is kept simple and straightforward. The scenario key frame transitions used as the basis for the Thumbwheel View prototypes were those that explored the user adjusting the view of the space from within, once he had placed himself in the space to see the results of furniture rearrangement.




Figure 6: Four frames from the Thumbwheel View prototype, representing the different views the user sees as he turns the roller wheel to rotate the view.

The Thumbwheel View prototypes for manipulating the user's view map the turning of the roller wheel to the rotation of the view left and right, as if the user were turning his head in the space. Moving the mouse on the desktop maps to moving the view position on the floor of the space, as if the user were walking across the floor. This mapping is consistent with the mapping of the thumbwheel mouse in the Thumbwheel Chair prototype, where turning the wheel rotates the chair on the floor and sliding the mouse forward and backward moves the chair across the floor. Informal user interaction with this interface proved successful; users thought of the interaction as smooth and natural.

The thumbwheel view prototypes also introduced additional interface features, such as graphical and sound feedback. When the user moves the mouse forward and backward to move the view forward and backward, the cursor changes to an arrow three-dimensionally formed to point in the direction of motion. Sound feedback is given as well, in the form of human footsteps. These footstep sounds help to simulate the notion of walking across the floor. Other visual feedback is given as well in the form of an additional window view of the space. This

additional view is included as an inset, to show an external view of the space, with the user's view represented by the location of a person in the space. As the user turns the thumbwheel roller wheel to turn the view of the space, the head of the person visible in the inset view window turns accordingly. Likewise, as the user moves the mouse to move forward and backward in space, the person in the inset view window moves in the space accordingly. While the graphical and sound feedback are included informally in these prototypes, they introduce a number of possibilities for interaction for future exploration. The mappings of the interaction with the thumbwheel mouse may need to be extended or customized, to allow for other interactions, as is the case in using the thumbwheel for object manipulation as well. However, as a start, the mappings seem to be appropriate and successful.

INTEGRATED APPLICATION PROTOTYPE

As a parallel effort to the MacroMind Director prototyping, Michael Chen incorporated many of the ideas from these prototypes, as well as additional interfaces, into a running application on the Mac IIx. The application is written in C++ using MacApp and a specially designed 3D graphics library. The building of an application allowed us to understand both the hardware and software requirements to actually use the interface techniques that were prototyped. Because it was able to integrate many of the potential interactive methods together in one application, it acts as an excellent platform to compare the different methods. The application also raised many overall interface issues that were not encountered in the more specific Director prototypes. For a more detailed description of the integrated MacApp prototype, refer to the MacApp Space Planning Technical Report forthcoming by Michael Chen.

Integrated Functionality and Constraints

The application uses simplified models of Swivel 3D furniture created for the Director prototypes. The simplified versions of the model were used to ensure responsive interaction speed in the program. While the Director prototypes were limited in their approach to the 3D manipulation tasks, isolating different types of interaction for prototyping purposes, the integrated application is a more dynamic general purpose manipulation program. For example, the application is able to provide for scale and perspective changes while the objects are sliding on the living floor. This was not possible in the Director prototypes.

The integrated prototype was also able to provide commonly expected default constraints for the objects in the living room. For example, the picture object is constrained to the wall on which it hangs. The user is free to place the picture at any point on the wall, or rotate it on the wall to adjust the angle at which it hangs. However, the user cannot inadvertently pluck the picture from the wall. Similarly, the furniture objects on the floor of the room are constrained to sit on the floor. The constraints are not meant to limit the user's interactions. They are there to guide the user so that the obvious interactions are made easy. Less prominent interactions may require more advanced tools. We recognize that in future iterations of the interface we need to provide the user with access to modify these constraints.




Figure 7: An overview of the MacApp Integrated Application, with each of the available manipulation methods in the leftmost column, and world and view manipulations on the borders of the window.

The different interactive methods currently available in the integrated application for manipulating 3D objects include:

1. Separate manipulation modes, namely slide and rotate. The user selects one of these and then performs the ensuing action on a selected 3D object.
2. Kitchen sink mode. The user performs slide or rotate actions on a selected 3D object, where the action performed is dependent on where on the object the user initially selected. The tool also combines the functionality of the Turn World tool (described below).
3. Handles. Similar to the kitchen sink mode, but incorporating the use of graphical handles from the earlier Director Handles prototype to give the user visual feedback as to the ensuing interaction.
4. Thumbwheel mouse. Allows the user access to the thumbwheel mouse to interact with objects as in the Thumbwheel Chair Director prototype,

where sliding the mouse maps to moving the object on the plane of the floor, and turning the roller wheel rotates the object around its axis.

The 3D manipulation interfaces are implemented using *virtual controllers*, developed by Michael Chen. Rotation is performed using the *virtual sphere*⁶. Other virtual controllers were developed to perform slide and lift functions.

World and View Manipulations

The integrated application also introduces manipulation techniques for maneuvering the 3D world itself, along with similar manipulations for controlling the user's view of the world. The methods explored by the integrated application do not currently incorporate all of the methods for manipulating views as introduced in the Director prototypes described earlier, instead presenting alternative means of interaction. One technique (the Turn World Tool) allows the world to be rotated around itself or in relation to a specific object in the scene using the virtual sphere. Another technique (the Turn About Tool) allows the world to be translated in 2D by allowing the user to "grab" and move the living room floor. Similar controls are available along the bottom and right panel of the application window for rotating and translating the world view. Also, zooming controls at the lower left corner of the window can be used to move the view inward and outward as desired.

User Testing and Iteration

As it stands, the integrated application is still not a completely functional 3D interface. It presents many of the interface elements necessary for a complete integrated interface, however in its current form it serves more as an interactive testbed for comparison among various techniques available. Many of the issues presented by this application, such as the role of constraints and which method to use for which interaction, still must be addressed and resolved to create a completed application. Further user testing is required to understand the tradeoffs among the various methods, and iteration of the prototype application is necessary to further refine the interface. However, the integrated application is currently an excellent testbed to resolve many of these issues because it presents so many of them simultaneously in a Macintosh environment, so the relationship of 3D to today's interface elements is considered. Further testing and iteration will help us to better understand the appropriate interfaces for 3D manipulation.

⁶The Convolvotron is a specialized hardware configuration developed at NASA which allows spatial placement of sounds in the audio field.

Incorporation of Localized Sound

As an experiment to explore the potential role of sound in 3D interaction, routines to play localized sound through the Convolvotron⁷ were incorporated in a version of the integrated application. By wearing headphones, the user can hear sounds associated with various interactions, such as the sound of a chair sliding across the floor as the user translates the chair. The sound appears at the corresponding 3D location with the object as it is moved around. While the role of sound must be further explored and tested to understand its appropriateness in interfaces for 3D manipulation, our initial work shows some success with using sound for interactive feedback. Audibly confirming the mode of manipulation, such as sliding a chair, can greatly enhance the realism of the interaction. It is our impression that sound which is related to the objects being manipulated in a specific, subjective way can help the user to understand complex interactions.

3D SOUND

Since its inception, the Leonardo Project has seemed to be a natural activity with which to integrate 3D graphical interaction and techniques, such as those that are evolving in the "virtual reality" centers of research such as NASA Ames. While we have decided not to pursue parallel research into the specialized helmet displays and multi-channel spatial input devices which are in vogue in that community, the localized sound technology which has been developed seems applicable as a useful cue for 3D manipulations by personal computer users. We thus began an exploration of localized sound as a supplement to the visual prototypes using the Convolvotron as noted in the previous section. The sounds are incorporated in the interface via stereo headphones, as are commonplace since the introduction of the Walkman™. We presented the idea of combining the interactive methods introduced in the Space Planning prototypes with localized sound cues and feedback. Other experiments in localized sound were conducted as well. The results of our exploration of localized sound will be summarized in the HIG Technical Report Sounds Good to Me, forthcoming by Lewis Knapp.

Our work touched on using sound as an environmental cue to support the illusion that the viewer is actually within the space. We prototyped a person moving from place to place in a scene which included the sound sources which were located acoustically with respect to the viewer's position. Our brief experiment indicated that this effect was subtle and not immediately evident to the casual user, but once understood the effect was noticeable. Perhaps with training which set users' expectations, this cue might be more valuable.

FINDINGS

The many approaches to 3D manipulation presented in the various prototypes are testimony to the fact that 3D manipulation is complex in nature. A single "best" approach is certainly not obvious, and may not be identifiable since different approaches seem appropriate for different types of tasks. While a single approach may not have been identified, however, overall findings can be summarized from the Space Planning project.

Interaction Based on Users' Expectations

The primary finding from all of the interactive prototypes is the identification of the need to map the 3D manipulation interaction to the user's real world expectations. The interaction must relate to the way the user thinks about the manipulation task in the physical world. For instance, in the physical world when the user moves a piece of furniture such as a chair, he receives feedback from the interaction such as the feel of the chair moving across the floor, the sound of the chair moving, and the responsive movement of the chair in relation to the user's movement. While each of these responses cannot be exactly recreated in a computer model, adequate feedback must be given to simulate the physical world. How much feedback is adequate depends on the interaction, but the user must "feel" that the interaction is accurate.

The prototypes developed seem successful based on this criteria of paralleling users' expectations of physical world interactions. The Handles prototype, for example, presents users with visual cues consistent with the way in which they might expect the object to be manipulable. Once one of the visual handles is selected, the ensuing manipulation consistently follows from their interaction. In the Thumbwheel Chair prototype, the chair is constrained to stay on the floor so that the user's interaction is limited to moving and rotating on the floor. Because for the most part these interactions are consistent with what users expect to do with the chair, the mapping of the sliding of the mouse and rotation of the roller wheel is understandable to most users.

Implications for Developers

The implication of this to developers of applications in other domains is that modeling the visual representation of the objects of interest is not enough; care must be taken to identify what properties of the objects in the domain of the application are most evident, obvious and important to the user. Simulation of these key properties is crucial to providing effective interaction with the objects.

In 3D domains like the Space Planning scenario, many of the properties are related to the shape of the objects, so that the internal computer representation used to generate the picture of the object can also be used to support interaction, eg. the picture slides on the wall. Providing the user with a recognizable display representation is not sufficient to support familiar, intuitive interaction with objects, and the developer must provide the additional element(s) which makes the interaction "work". By constraining objects, we simulated properties such as "sits on the floor", "rotates (only) about the vertical axis" and "can be lifted". These properties can be expressed in geometric terms which could be provided as part of a modeling toolbox so that developers would have a common base for their simulations. More work is needed to identify and generalize these aspects which support simulation for interaction.

Manipulating People v. Objects

We also found that users have different expectations for manipulating people than they do for manipulating inanimate objects. Because most users do not directly manipulate people in the physical world, interactions for maneuvering people seem more indirect or delegatory in nature. The smooth integration of a delegation-oriented interface with support for direct manipulation where appropriate is introduced here, but continued exploration is necessary.

Flexible Interchange of Methods

We expect that individual users will have preferences for different interactive methods dependent on the task undertaken. Any of the methods we have prototyped have sufficient generality to manipulate 3D objects which are different from chairs and tables, and could form the basis of a tool set which could be used consistently across different domains. We believe that specific applications may require more specialized methods depending on the domain, but expect that they would be used in addition to a generally understood core of interactions like the ones we have demonstrated.

An architecture which allows the developer and end user to interchange and customize these interaction methods is desirable to permit a user to select the methods he prefers. Similarly, such an architecture would allow the interactive method to change based on the hardware available (for instance if a thumbwheel mouse was accessible or not). At the very least, application developers must recognize that each interactive method may not be appropriate for every 3D manipulation task, and should employ methods that are most consistent with users' intentions.

Role of User Testing

These findings are based on our experience developing the 3D interaction methodologies, and informal discussions with users of these methods. We recognize the need for more formal user testing to identify concrete findings and recommendations. Further development of these ideas and methods will require and benefit from more extensive user testing of the techniques in a setting which integrates several of them so that comparative tests are possible.

3D Prototyping Isn't Easy!

Finally, through developing these prototypes, we learned how difficult 3D prototyping is on the Macintosh. Very few 3D manipulation tools exist, and those that do, such as Super3D and Swivel 3D, are primarily modeling tools. They do not provide alternatives for interacting with the 3D objects created in these applications. Thus, to create interactions with 3D objects, we first had to model the objects in one of these environments, then bring them in two-dimensional pieces into an interface prototyping environment (we used MacroMind Director). We understand that 3D Works, a 3D version of Director soon to be released by MacroMind, would alleviate the task somewhat, but not completely. Interaction methods with 3D objects are built into each of the manipulation applications. To introduce new methods for interaction, we basically have to make our own applications (as we did with the MacApp Integrated Application). While the tools we had were limited, we found comfort in knowing that a greater range of prototyping tools seems to be available on the Macintosh than on any other platform, though we did not extensively explore the possibility of prototyping on other machines.

ISSUES

A number of issues related to 3D direct manipulation interaction have been identified over the course of development of the Space Planning project. These issues should be considered in building 3D interfaces, and are summarized below.

Input Device

Knowing the input device before developing the interface is important to building an interactive dialog between the user and the objects being manipulated. By identifying the input device before developing the interface, interface designers can be more sure of a consistent mapping between the user's expectations of the way the input device will work and the interaction with the 3D objects. Interfaces for different devices might indeed be quite different.

Device mappings and user customizability

In addition to mapping the use of the input device to user expectations, other concerns about the nature of the device and the interaction must be considered. Specifically, because the devices used for 3D manipulation typically have less than the six degrees of freedom necessary to map user manipulation of the device directly to manipulation in the computer-modeled environment, the mappings of the device to the interface will only allow manipulation of a subset of these degrees of freedom. Even when input devices have additional degrees of freedom, the primary output device, the computer monitor, has only two dimensions, so mapping the interaction accurately to the screen is also a problem. Allowing the user to specify which degrees of freedom are controllable by the input device, or customize the constraints of the interface in general, is desirable.

Integrated Interface

The prototypes discussed above typically isolate different interactions for exploratory purposes. Even the integrated MacApp application, which combines many of the ideas explored, does not finalize the integration of the interface to provide a single interface for 3D manipulation. Future research remains to be done to identify the elements of an interface desirable and necessary for a successful integrated interface. How these tools would be made available in a component based or application based architecture is also unexplored at this time. The approach to potentially multiple views onto the 3D space must also be resolved.

Role of sound

While sound as feedback was explored in some of the prototypes and the integrated application, and in parallel experiments with the Convolvotron, the potential for sound in 3D interaction has yet to be fully developed and tested. We believe that the use of sound can be important to the success of the interface, both as feedback and potentially as interface cues. Using sound as interface cues and/or feedback may allow the interface to compensate for other missing modalities of feedback, such as tactility. The exact role of sound and its relative importance in 3D interaction is yet to be determined.

Object Constraints and Default Behaviors

The importance of constraints on object manipulations, both to help users complete desired tasks and to simplify the interaction, has been introduced in the Director prototypes and the integrated application. Object constraints might be user-adjustable, or might change based on the interactive method being applied, but the most obvious conclusion is that they must be provided to the user in a way which meets his expectations to begin with. Further exploration into the role of constraints and the effects on the overall interface must still be continued. Related to object constraints, objects might be given default behaviors which users can then adjust. Desirable default behavior, based on consistent user expectations, is yet to be determined.

CONCLUSIONS

Overall, the Space Planning project demonstrates that a number of successful solutions to the problem of 3D direct manipulation are possible. The variety of methods demonstrated in the project stresses that different 3D approaches will make 3D manipulation accessible to most users. Smooth and consistent interaction is most definitely possible in 3D interfaces.

While no single "best" solution to the problem of 3D manipulation has been identified, and further user testing remains to be done to draw conclusive results, the project as a whole was successful. With the Space Planning project, we have finally provided an appropriate interactive scenario for 3D manipulation, and a number of valid approaches to 3D interaction. We have also identified key issues corresponding to 3D direct manipulation interaction. Manipulating objects in 3-space is a desirable and achievable task, and interfaces to do so provide users with new functionality for their personal computers. Our indication that 3D manipulation is attainable and beneficial for the user is the first step to further advances in intuitive 3D interface.